New Mission, Old Spacecraft: EPOXI's Approach to the Comet Hartley-2

Richard R. Rieber* and Gregory R. LaBorde*

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109

NASA's Deep Impact mission ended successfully in 2005 after an impact and close flyby of the comet 9P/Tempel-1. The Flyby spacecraft was placed in hibernation and was left to orbit the sun. In 2007, engineers at the Jet Propulsion Laboratory brought the spacecraft out of hibernation and successfully performed two additional missions. These missions were EPOCh, Extra-solar Planetary Observation and Characterization, a photometric investigation of transiting exo-planets, and DIXI, Deep Impact eXtended Investigation, which maneuvered the Flyby spacecraft towards a close encounter with the comet 103P/Hartley-2 on 4 November 2010. The names of these two scientific investigations combine to form the overarching mission's name, EPOXI.

The encounter with 103P/Hartley-2 was vastly different from the prime mission's encounter with 9P/Tempel-1. The geometry of encounter was nearly 180° different and 103P/Hartley-2 was approximately one-quarter the size of 9P/Tempel-1. Mission operations for the comet flyby were broken into three phases: *a*) Approach, *b*) Encounter, and *c*) Departure. This paper will focus on the approach phase of the comet encounter. It will discuss the strategies used to decrease both cost and risk while maximizing science return and some of the challenges experienced during operations.

Nomenclature

DDOR	Delta-Differential One-way Ranging
DIXI	Deep Impact eXtended Investigation
DOY	Day of Year
DSN	Deep Space Network
EPOCh	Extra-solar Planetary Observation and CHaracterization
HGA	high-gain antenna
HRI	high-resolution imager
MRI	medium-resolution imager
TCM	trajectory correction maneuver
TWTA	Traveling wave tube amplifier
UTC	Universal Time Coordinated
VTC	Vehicle time code

I. Introduction

NASA's Deep Impact mission, launched 12 January 2005, was a smashing success after its impact and close flyby of the comet 9P/Tempel-1 on 4 July 2005.[1] A few weeks later, the spacecraft had served its purpose and was placed in hibernation and left to orbit the sun. In 2007, engineers at the Jet Propulsion Laboratory with support from the spacecraft manufacturer, Ball Aerospace Technologies Corporation, brought the spacecraft out of hibernation to perform two additional scientific observations. The primary objective was DIXI, or **D**eep **I**mpact e**X**tended **I**nvestigation which maneuvered the spacecraft towards a close flyby of the comet

*Staff Engineer, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91109

103P/Hartley-2 on 4 November 2010.[2] During the first year of the 3-year cruise to Hartley-2, the flyby spacecraft conducted a secondary scientific objective entitled EPOCh, or **E**xtra-solar **P**lanetary **O**bservation and **Ch**aracterization, a photometric investigation of transiting exo-planets[3,4] and observing Earth as an exoplanet analog.[5] The two acronyms identifying the scientific campaigns, EPOCh and DIXI, together form the mission's new name: EPOXI.

NASA's spacecraft are traditionally optimized to satisfy specific scientific objectives. For example, the geometry of the Deep Impact Flyby spacecraft is optimized for Tempel-1 approach: the cameras were pointed at Tempel-1 while the high-gain antenna (HGA) was pointed at Earth during Approach; key components were shielded from damage caused by dust in the comet's coma on that specific trajectory; and the instruments were shaded by the solar panel to keep them cool at the expected attitude. When NASA missions are granted an extended mission, for example Cassini or the Mars Exploration Rovers, the spacecraft and the mission operations team continue their existing campaigns and observations, well-suited to their designs.

What makes the EPOXI mission unique is that it is not the extended mission of Deep Impact. Rather it is a mission of opportunity that makes use of the existing Deep Impact Flyby spacecraft, ground support systems, and a few key team members to perform two completely different missions. The mission operations of the EPOCh portion of the mission have already been discussed in reference 3. This report will focus on mission operations around the Hartley-2 flyby with an emphasis on 60-days from closest approach until 1-day from closest approach, called the Approach phase.

I.A. Flyby Spacecraft

The Deep Impact Flyby spacecraft, seen in Figure 1, is 3-axis controlled and is equipped with two telescopes: the high-resolution imager (HRI) and the medium-resolution imager (MRI). The HRI is equipped with both a visible camera and an infrared spectrometer whereas the MRI is only equipped with a visible camera. Both cameras are furnished with nine filters. The field of view of the HRI and MRI cameras are 2-mrad (0.118°) and 10-mrad (0.587°) respectively. The slit of the IR spectrometer has a field of view of 10-μrad by 2.5-mrad (0.00057° by 0.15°). There is no scan mirror in the spectrometer, thus, the entire spacecraft must slew to scan the slit across the desired target. More details about the instruments can be found in reference 6.

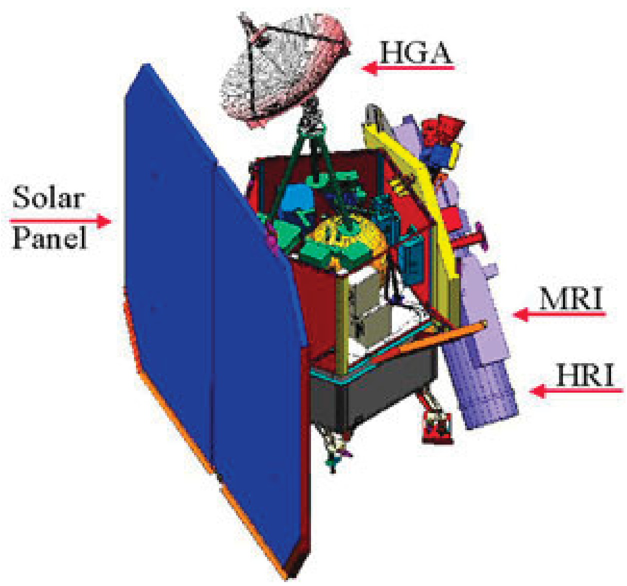

Figure 1. The Flyby spacecraft with the top deck and one side panel removed.

The design of the Flyby spacecraft was optimized for the Tempel-1 encounter. The instruments were positioned to view Tempel-1 on Approach while pointing the HGA at Earth. Furthermore, the spacecraft was equipped with Whipple shields in strategic locations based on the flyby geometry to protect critical components from dust impact when the spacecraft flew through Tempel-1's coma. The geometry of the encounter with Hartley-2 was completely different. When the spacecraft was imaging the comet on Approach, the HGA was pointed about 140° from Earth. This rendered communications while imaging impossible during Approach. Thus, the flight team had to adopt a completely new strategy for the Hartley-2 encounter.

The design of the Tempel-1 encounter centered around the idea of "live for the moment." This strategy attempted to preserve the spacecraft computer, telecommunications equipment, and instruments for real-time data collection and downlink in the event the spacecraft was incapacitated in the next moment. Significantly less effort was expended to protect components not critical to those tasks. The Hartley-2 encounter design centered around the idea of "live for the playback," as data could not be downlinked until after closest approach and had to be saved to on-board memory for a later playback. This design strove to preserve many other aspects of the spacecraft as many more components were required for a future playback.

I.B. Encounter Overview

The flyby of Hartley-2 was an affair that lasted 82-days. Throughout the entire Encounter, times are referred to relative to their proximity to of the time of closest approach, also called Encounter. Thus, seven days before encounter would be E-7d and similarly, seven days after encounter would be E+7d. The 82-days of the Hartley-2 encounter were broken up into three phases summarized in the following subsections. Each phase was characterized by unique science requirements and engineering constraints that will be briefly discussed.

I.B.1. Approach

Approach encompassed E-60-days until E-26-hours. Approach was broken up into several sub-phases that will be discussed in detail throughout this paper. Although the scientific goal of Approach was continuous data collection of the comet, it was punctuated with calibrations and three trajectory correction maneuvers (TCM). This was made even more challenging by the geometry of the Flyby spacecraft and the locations of the Earth and Hartley-2. The orbital geometry made communication via the HGA impossible while imaging the comet and necessitating a large slew away from comet-pointing to downlink data.

I.B.2. Encounter

Encounter lasted from E-18-hours until E+2-days after closest approach. The spacecraft began by continuously imaging the comet from the beginning of the phase through closest-approach at a distance of 694 km and a closing velocity of 12.32 km/s. The spacecraft was unable to downlink any data before closest approach owing to the geometry of the encounter, thus all data was saved to onboard storage. After closest approach the geometry allowed the HGA to point at Earth so the spacecraft downlinked as much data as possible while continuously imaging the comet.

I.B.3. Departure

Departure began when the Encounter sequences ended at E+2-days. Imaging continued until E+21-days. The geometry of Departure was such that the spacecraft could image the comet and point the HGA at Earth simultaneously. The challenge was to collect data no faster than it could be downlinked.

I.C. Approach Organization

Approach was divided into three sub-phases Early-Approach, Mid-Approach, and Late-Approach, each with its own science, engineering, and operational requirements. These sub-phases were bounded by periods of time when no science data was collected. It was during these interludes where a variety of activities were accomplished such as instrument calibrations and TCMs. Due to thermal constraints discovered late in planning, Early-Approach was further sub-divided into two sub-phases. The exact schedule can be seen in Table 1. Each phase will be discussed in detail later in the paper.

II. Constraints

Compared to other Discovery-class missions, EPOXI was run on a shoestring budget. It utilized an existing spacecraft and an existing mission operations infrastructure established almost a decade ago to perform a completely different mission. Furthermore, the original EPOXI proposal submitted to NASA stated that there would be no flight software modifications.

In late 2007, the original target for the EPOXI mission, the comet 85P/Boethin, could not be found. Unfavorable orbital geometry during the comet's previous perihelion pass rendered it non-observable from Earth.[7] An extensive ground-based observation campaign could not detect the comet. It was decided to re-target EPOXI towards the comet 103P/Hartley-2.[8] This decision was not made lightly as the retargeting increased the mission duration from 16-months to 41-months, over 250% longer, with no increase in budget.

Over the prime mission and the EPOCh portion of EPOXI, the Flyby spacecraft has proven itself to be finicky to operate. JPL and NASA management came to the conclusion that the largest risk to the EPOXI mission would be to hibernate the spacecraft, disband the flight team, and bet they would return to wake up the spacecraft and fly it past the comet. They decided to keep the flight team fully staffed for the entire duration of the cruise to 103P/Hartley-2. This decision necessitated keeping the team small and agile to save

Table 1. Approach Schedule

Activity	Relative start	UTC start
Early-Approach	**E-60d**	**2010-248T13:12:19**
Early-Approach #1	E-60d	2010-248T13:12:19
Early-Approach #2	E-50d	2010-258T15:03:57
First interlude	**E-40d**	**2010-268T16:01:49**
Instrument cool-down	E-40d	2010-268T16:01:49
Instrument calibration	E-37d	2010-271T21:38:48
TCM-20	E-36d	2010-272T16:57:01
Mid-Approach	**E-34d**	**2010-274T15:59:57**
Second interlude	**E-8d**	**2010-300T16:58:00**
TCM-21	E-8d	2010-300T16:58:00
Late-Approach	**E-8d**	**2010-300T22:44:59**
TCM-22 (contingency)	E-2d	2010-306T13:58:00
Late-Approach (continued)	E-2d	2010-306T17:46:58
Transition to Encounter	**E-26h**	**2010-307T11:59:58**

cost. No extra staff was hired to improve or modify infrastructure established for the Deep Impact prime mission.

In addition to the optimization of the Flyby spacecraft for the Tempel-1 encounter, the spacecraft and ground software were never designed to last longer than about a year. The Hartley-2 encounter completed when the Flyby spacecraft was nearly 6-years old, over six-times its design life. As such, some of the hardware began to show its age. The main subsystem that was causing operational issues was the downlink path of the telecommunications equipment. The Flyby spacecraft is equipped with redundant downlink paths. The heart of each downlink path is the Traveling-Wave Tube Amplifier (TWTA). Each TWTA is connected to an antenna via a waveguide transfer switch in such a manner that one TWTA is always connected to the low-gain antenna (LGA) and one is always connected to the HGA with only a single TWTA powered at any given time. To swap between the LGA and HGA, mission operators simply toggle the waveguide transfer switch. However, if the switch cannot be toggled, one can swap antennas by powering off one TWTA and power on the other.

In February 2008, the prime TWTA (TWTA-A) began showing symptoms that it may be aging. Project management decided to swap to the backup TWTA (TWTA-B) to preserve the life of TWTA-A. This TWTA situation was further exacerbated in April 2008 when signal strength (from TWTA-B) lost 8dB (75%) of its power. This is briefly mentioned in reference 3. The leading hypothesis of this anomaly was a thermal issue created because the spacecraft was closer to the sun than it had ever been before causing higher temperatures. As the spacecraft progressed in its orbit and the spacecraft–sun distance increased, the signal strength was restored to its full power lending credibility to this hypothesis. However, there was a competing hypothesis of debris in the waveguide. If this were to be true and the debris worked its way into the waveguide transfer switch, there was a risk that the waveguide transfer switch could become stuck between its two positions when toggled. This would block any radio signal from getting out of either antenna and end the mission. This is similar to what was experienced on the Mars Reconnaissance Orbiter,[9] fortunately its switch is stuck in one of its operational positions.

The possibility of a mission-ending failure necessitated a ban on toggling the waveguide transfer switch. The only method now available to the flight team to swap between antennas is to power off one TWTA and power on the other. Routine power-cycling of TWTAs was not recommended by the manufacturer and the flight team made the decision to drop communication when turing off-Earth for science imaging. This was comically referred to as the "no-gain antenna."

III. Sequence Structure

All commands for Approach were built into sequences. Very little commanding was done in real-time. A sequence is simply a time-ordered list of commands and can exist in two flavors: relative-timed or absolute-timed. A relative-timed sequence begins clocking out its commands as soon as it is activated. The timing of the execution of a command is relative to its neighboring commands. For example, execute command 1, wait 30-seconds, execute command 2. An absolute-timed sequence executes its commands at a specific time. The flight team would think of each command executing at a specific UTC, however, the execution time of commands in an absolute-timed sequences are specified in vehicle-time code (VTC) and not UTC. A complex procedure of mapping VTC counts to J2000 UTC helps predict what the VTC will be at a specific UTC.

Relative-timed sequences were used in Approach to create utility sequences that could be executed multiple times. For example, the commands configuring the telecommunication subsystem for image downlink was built into a stand-alone relative-timed sequence. It was executed each time the spacecraft returned from imaging to downlink data. In the event the flight team desired a different telecommunication configuration for data downlink, only a single sequence had to be modified. This utility-sequence strategy was used for turns routine slews, telecommunication configuration, and a plethora of other situations. In addition to utilities, all imaging commands were built into camera-specific relative-timed sequences. Thus, imaging commands were executed relative to the time the sequence was called and allowed the sequence to be utilized multiple times.

The only absolute-timed sequences used for Approach were what the flight team refers to as a "backbone sequence." The purpose of the backbone sequence is to call relative-timed sequences at an appropriate time and to synchronize on-board operations with activities on the ground, such as a DSN contact. Thus, each phase of Approach included a single, absolute-timed backbone, a relative-timed sequence for each of the three detectors (MRI, HRI-visible and HRI-infrared), a relative-timed sequence for all of the slews to and from the comet (referred to as the "engineering sequence"), and a series of utility sequences used throughout all of Approach.

IV. Early-Approach

Early-Approach began on DOY 248 at E-60d. The challenge to this early phase of Approach was the imaging attitude causing thermal issues. While imaging the comet, the top-deck and instrument platform of the spacecraft was in the sun as seen in Figures 2(a), 2(b), and 2(c). The solar illumination had two main effects on operations: *a*) the infrared detector warmed up and made the collection of infrared spectra impossible and *b*) the telecommunication components, primarily mounted to the top-deck, warmed up, notably, the problematic TWTAs, increasing the risk of the 8dB loss returning.

Science recognized extremely early in planning that infrared data would be impossible to collect during Early-Approach because of sunlight warming the infrared detector. However, the thermal effects on the telecommunication components mounted on the underside of the top-deck of the spacecraft (seen peeking out from behind the solar array shown in green in Figure 2(a)), were discovered much later in the development of the Early-Approach sequences. These effects drove the need to split Early-Approach into two 10-day activities. Each activity with a similar structure, only a different science collection cadence to minimize solar heating.

The science team requested data sets every two-hours during Early Approach. Each data set consisted of nine HRI images and twelve MRI images. In addition to the science data, the optical navigation team requested a single MRI image taken with an exposure duration selected to reduce signal from the coma and maximize signal from the nucleus, but still retain background stars in the field of view.[10] However, because of the thermal issues the sampling rate for the first ten-days of Approach had to be reduced to once every six-hours to allow the top-deck to cool.

For this phase of approach, there was a single absolute-timed backbone and four relative-timed sequences, one each for the MRI, the HRI, the routine engineering activities, and one for data downlink. All four sequences were constructed to collect and downlink twelve data sets. This strategy let each data set reside in the spacecraft's memory for a long duration of time providing multiple chances for it to be downlinked.

Each cycle of data collection began by reconfiguring the telecommunication subsystem to use the "no-gain" antenna, described previously, then continued with a turn to the imaging attitude. Once at the imaging attitude, the spacecraft remained stationary for 10-minutes to allow attitude errors to damp out, then began

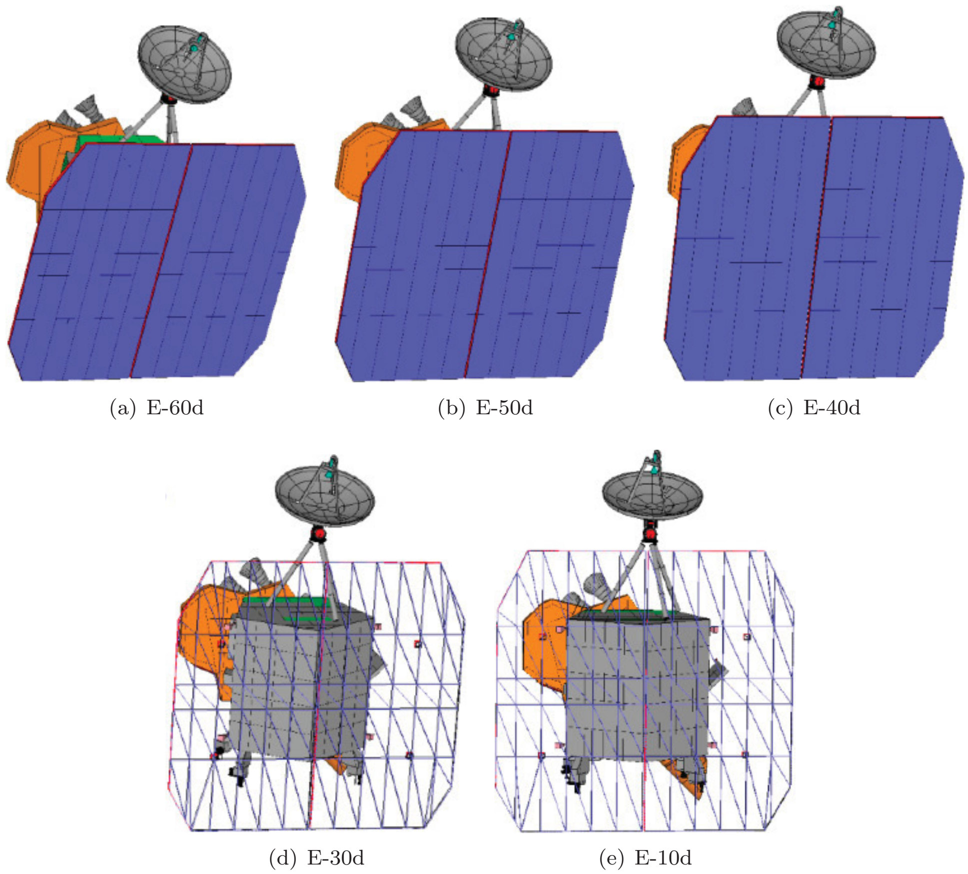

Figure 2. Imaging attitude at various phases during Approach as seen from the sun. Note the sunlit top-deck in green at in (a). Image courtesy of thermal engineer John Valdez, Ball Aerospace Technologies.

imaging. The imaging commands were built into camera-specific sequences such that images from the HRI and MRI could be taken simultaneously. When imaging had completed, the spacecraft turned back to a playback attitude. Note that this is not the routine cruise attitude used for the past three years, but a skewed attitude that decreased the angle between the imaging and playback attitudes while still allowing the HGA to point at Earth. Once at the playback attitude, data could begin to be downlinked to the DSN.

IV.A. E-60 to E-50

Even with a 6-hour imaging cadence when only 00:34:05 was spent imaging the comet, the thermal environment remained inhospitable to the telecommunication components. Reducing the slew durations from 600-seconds to 375-seconds provided just enough extra time to bring the telecommunication components back under their limits. This imaging cadence and rapid slews left 05:12:25 out of every 6-hours with the HGA pointed at Earth. This afforded plenty of time to downlink twelve-cycles, or three-days worth of data. Every time the HGA was pointed at Earth, the spacecraft would downlink every data set twice regardless of whether there was a scheduled DSN track. This decoupled the DSN schedule from spacecraft operations and allowed these sequences to be built and tested well before a DSN schedule was negotiated and delivered.

During EPOCh observations, described in reference 3, the 8dB loss occurred and the data rate had to be cut in half causing data loss even though every image was transmitted twice. Because the leading hypothesis for the cause of the 8dB loss was increased temperatures of the telecommunication components and the imaging attitude placed the spacecraft's top-deck in the sun, the design had to allow for the return of the 8dB loss. If the spacecraft were to lose 8dB of its signal strength, the data rate would have to be cut in half. The downlink was structured such that for one cycle, the playbacks were ordered newest to oldest. For the next cycle, the playbacks were ordered from oldest to newest. This flip-flop ensured that all data would be downlinked twice at the nominal data rate of 400kbps, but would also be downlinked once at the reduced data-rate of 200kbps. This technique works because of a complex downlink queue management, that was also used for EPOCh imaging, described in reference 3.

In a further attempt to reduce the risk of telecommunication components warming up, the temperature of TWTA-B was closely monitored during Early Approach. The flight team came up with a simple contingency plan: if the temperature of TWTA-B reached 53°C, the HRI would be powered off to eliminate a heat source inside the spacecraft bus. After a day's worth of imaging and downlinking, the TWTA temperature had in fact, reached the 53°C limit and this contingency procedure was carried out. The TWTA temperature dropped by 2.5°C to a safe value. No further HRI data was collected until it was powered back on.

Because of the massive amounts of downlink redundancy, only one 6-hour DSN track was needed every three days to ensure all data was received. However, this was not the case and EPOXI was given two 6-hour DSN tracks each day ensuring that the science team was overwhelmed with sometimes as many as 8 copies of every image.

IV.B. E-50 to E-40

The cadence of the E-50 to E-40 sub-phase of Approach reduced from once every six-hours to once every 2-hours. The duration of a complete cycle reduced from three-days to 24-hours. The slew durations were also relaxed from 375-seconds to 600-seconds. The imaging sequences remained the same, albeit with a shorter cadence.

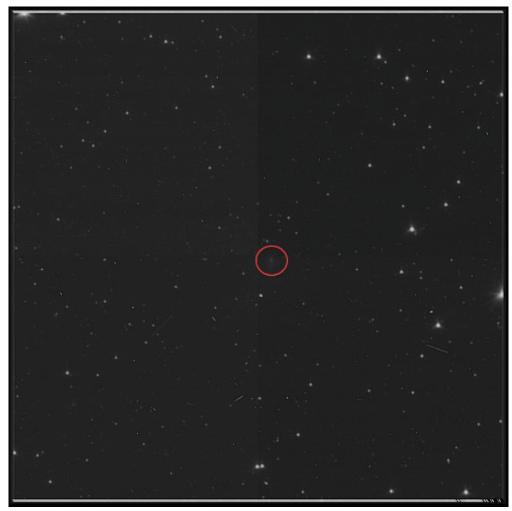

Figure 3. The first picture of Hartley-2, inside the red circle, taken at E-60d.

The playback sequence had to be completely revamped. Only 01:04:55 out of every two-hours was spent with the HGA pointed at Earth. It took 00:53:29.5 to downlink 12 data sets at the nominal 400kbps. 00:10:15 was also required to downlink logs, directory listings, and engineering health and status data. This summed to a total downlink time of 01:03:44.5, which theoretically fit within the 01:04:55 spent Earth-pointed. The only concern is when the spacecraft was imaging, no radio signal would reach Earth. When the spacecraft slewed back to the playback attitude and pointed the HGA at Earth, this would be the first signal the DSN saw from the spacecraft and needed a bit of time to lock onto the signal. Only two-minutes was allocated to the DSN to lock onto the signal and appropriately configure for data downlink. This meant that a complete downlink of 12-data sets, logs, directory listings, and engineering data required 00:00:49.5 more time than was available, necessitating two cycles to receive every image on the ground once.

If the 8dB loss were to be experienced again, the data rate would have to be reduced from 400kbps to 200kbps. The playback sequence had to be designed to operate at both data rates. The image sets were placed in the downlink queue in order of newest to oldest for one cycle, then oldest to newest for the next cycle. This maximized the number of unique images being received on the ground if the spacecraft had to be commanded to a lower data rate.

On DOY 263, the HRI was powered back on. The thermal environment was now much less hostile to the telecommunication componetns and data collection with the HRI could continue.

Fortunately, throughout Early-Approach, the 8dB loss never returned. The orbit prediction of the comet was fantastic, leading to excellent spacecraft's pointing. The first picture taken of Hartley-2 from the Flyby spacecraft can be seen in Figure 3

V. First Interlude

Early-Approach ended on DOY 268 and returned from the skewed playback attitude to cruise attitude with the solar array normal to the sun. The spacecraft sat in this attitude for three-days, until DOY 271. At this attitude, the entire spacecraft bus and instrument platform were in the shadow of the solar array, keeping all components cool, but most importantly, allowing the infrared detector to reach its operating temperature. With the infrared detector cool, the spacecraft executed a day-long calibration activity that had been used several times throughout the mission.

On DOY 272, the spacecraft executed trajectory correction maneuver number 20 (TCM-20). The challenge with flying past a comet is its orbit is difficult to predict. As a comet flies through space, it also rotates. When certain portions of the comet's surface is illuminated by the sun, it heats up and frozen water begins to vaporize, creating jets of material. This is best seen in a Figure 4, taken during the flyby. Notice the bright, radiating lines extending from the bottom right of the comet. These are jets of water and dust. These jets act like small, periodic thrusters and change the orbit of the comet. TCM-20 corrected the orbit of the Flyby spacecraft based on the optical navigation pictures taken during Early-Approach. After the completion of TCM-20, the Flyby spacecraft returned to collecting data in Mid-Approach.

VI. Mid-Approach

Figure 4. Jets coming off the surface of Hartley-2 that change its orbit, making navigation difficult. This image was taken during the flyby and not during Approach.

Mid-Approach began at E-34d on DOY 272. This phase of Approach was perhaps the most relaxed. It lasted for 26-days and was characterized by 16.5-hours of imaging the comet and 7.5-hours with the HGA pointed at Earth. As with Early-Approach, Mid-Approach was structured with a single, absolute-timed backbone sequence and five relative-timed sequences, one each for the three detectors, the engineering sequence, and a playback sequence. A data collection cycle began when the backbone called the four sub-sequences. These sequences occupied the spacecraft for 16.5-hours and concluded with the spacecraft pointed back at Earth. Then, based on the DSN schedule, the playback sequence was called. Downlinking all the data took about 5.5-hours, leaving two hours unoccupied.

315MB of data was collected each day. Of which, 43% was HRI-IR, 31% was MRI, and 23% was HRI-Visible. The remaining 3% of data was optical navigation data. This was taken with the MRI at a one-hour cadence for a total of total of 16 navigation images.

The two spare hours at Earth-point proved to be vital. It let the start of the playback sequence shift earlier or later, and even let the start of the sub-sequences shift earlier or later. Most playbacks occurred over the Goldstone DSN complex in California. By shifting the playback later, it left time at the start of the Earth-pointed window where the spacecraft could conduct advanced ranging activities called a Delta-Differential One-way Ranging (DDOR). DDORs require the spacecraft to be in view of two DSN complexes. Shifting the playback later allowed for DDORs between Goldstone and Madrid DSN complexes. Similarly, shifting the playback earlier allowed for DDORs between Goldstone and Canberra DSN complexes.

This extra time at the playback attitude also gave the flight team an opportunity to capture time correlation packets from the backup flight computer. Each flight computer on the Flyby spacecraft has its own clock. These packets help map the computer's VTC to UTC. The telecommunication system is configured such that the DSN can listen to only a single flight computer at any given time. The clock on the primary flight computer was well understood since the vast majority of time is spent listening to it. Absolute-timed sequences require an accurate understanding of the computer's clock. Without this understanding, there is no guarantee at what UTC commands will clock out. In the event the primary flight computer failed during the encounter imaging, the backup flight computer would take over and begin running a similar absolute-timed sequence to that executing on the primary flight computer. The only difference between the primary and backup computer's absolute-timed sequences is the VTC-UTC correlation. This correlation is generated from the time correlation packets.

During Mid-Approach, the extra time at the playback attitude allowed the flight team to switch the telemetry source to the backup flight computer, generate ten-minutes of time correlation packets, and swap back to the primary flight computer. This entire activity took 13.5-minutes plus round-trip light-time between Earth and the spacecraft, about 160-seconds. This mini-activity involved two relative-timed sequences and a single real-time command. One relative-timed sequence was installed on the primary flight computer. It switched the telemetry source from the primary to the backup computers, then, 13.5-minutes later, swapped the telemetry source back to the primary computer. The second relative-timed sequence was installed on the backup computer. It simply generated time correlation packets for ten-minutes and concluded with a downlink of the command log. The real-time command actually included two commands. The first activates the primary computer's sequence and the second activates the backup computer's sequence. The mission controllers were given the authority to send this command without requiring additional approval. They were also provided a table of time windows when this command could be sent. This technique facilitated the most accurate mapping of the backup flight computer's clock since launch.

VII. Second Interlude

Mid-Approach concluded at 2010-300T16:58:00 and left the spacecraft at the nominal cruise attitude in the nominal cruise configuration. The only activity during the second interlude was TCM-21. The TCM-21 sequence picked up operations from this cruise configuration. The direction and magnitude of the burn was calculated based on the optical navigation data taken during Mid-Approach. TCM-21 executed flawlessly and returned the spacecraft back to a nominal cruise configuration in preparation for Late-Approach imaging at 2010-300T22:44:59.

VIII. Late-Approach

Late-Approach was perhaps, the most stressful portion of EPOXI mission operations to date. Both the navigation team and science team required hourly sampling of the comet. With the Flyby spacecraft unable to simultaneously image the comet and the point the HGA at Earth, meeting these requirements was exceedingly difficult, however, the flight team was able to meet them with a complex series of maneuvers, we jokingly began to refer to as the do-si-do, after the American square dancing step.

The science team's baseline request was a 17-hour session of continuous imaging of the comet, followed by 7-hours of the do-si-do. The do-si-do was a series of maneuvers that slewed the spacecraft between comet-point and Earth-point each hour. The spacecraft began at comet-point, slewed to Earth-point in 00:06:15, downlinked as much data as possible, then returned to comet-point with a 00:06:15 slew to continue imaging. A total of seven hourly do-si-dos were performed each day.

Late-Approach began immediately after TCM-21, at 2010-300T23:00:00 and ended at 2010-307T12:00:00. This phase encompasses 6-days, 13-hours. The science team's baseline request encompassed 24-hours of data collection. Somewhere within this 6-day, 13-hour window, something had to flex. The 24-hours of science collection was broken up into two portions, 17-hours of comet-staring and the 7-hour do-si-do. Each portion had four relative-timed sequences, three each for the detectors and the engineering sequence. The do-si-do sequences were able to halt the comet-staring sequences and gave the flight team the flexibility needed to gather data for the entire 6-days, 13-hours allocated to Late-Approach. Both groups of sub-sequences were called from a single, absolute-timed backbone sequence.

VIII.A. Comet-Staring Sequences

The comet-staring sequences were rather basic. The engineering sequence began by reconfiguring the telecommunication subsystem for imaging, then slewed the spacecraft in 00:06:15 to point the instruments at the comet. From there, it conducted an hourly schedule of IR scans and comet-pointing. The final command in the engineering sequence was an IR scan. Had this scan not been halted, the spacecraft would continue its scan until it hit a fault limit. The scan is only scheduled to last for 352-seconds. The engineering sequence assumes that the do-si-do sequences will pick up and stop the scan to continue imaging. However, if for some reason, the do-si-do sequences do not begin, the comet-staring engineering sequence had a fail-safe command to slew back to the playback attitude. Theoretically, the engineering sequence would be killed by the do-si-do engineering sequence before this command was ever allowed to clock out.

Each detector sub-sequence executed a series of imaging commands with varying filter wheel positions and exposure durations. The details of which can be read about in 2 and its references.

VIII.B. Do-Si-Do Sequences

The do-si-do sequences picked up where the comet-staring sequences left off. They begin assuming the spacecraft is already pointing at the comet and appropriately configured for imaging. The sequences immediately begin taking data for about 12-minutes, after which the spacecraft begins the do-si-do maneuvers. The spacecraft swings around rapidly and as soon as it is Earth-point, the telecommunication subsystem is configured for image playback. The challenge with sequencing the start of image transmission is to predict how long it will take the DSN to lock on to the spacecraft's signal and appropriately configure the station. Traditionally, the station turns on its transmitter about 5-minutes after start of track and sweeps the uplink frequency. Once the spacecraft's receiver has locked onto the uplink sweep, it generates a more accurate downlink frequency. When this downlink is received by the DSN, communications is considered coherent and is referred to as in-lock, 2-way. The uplink sweep takes 2-minutes and a coherent signal is received a round-trip light-time later. For this portion of Approach, the round-trip light-time was about 158-seconds. Once the station was 2-way, it would then take about another 5-minutes to enable ranging and configure a few additional parameters. Thus, the station was ready to receive data about 15-minutes after the start of a track.

For the do-si-do, this was not acceptable. The downlink of the data collected would take 93% of the time when the spacecraft was Earth-pointed. This is 195-minutes of downlink out of the 211-minutes of Earth-pointed time. A 15-minute lockup each do-si-do contact would total 105-minutes, which is about half of the time the spacecraft is pointed at Earth.

The solution was to appropriately configure the station prior to the spacecraft pointing at Earth. Remember, these activities are performed using the "no-gain" antenna and not even carrier signal is seen at Earth, let alone telemetry. The maximum LGA boresight to Earth angle was 66° throughout all of Late-Approach. Because of the broad beam-width of the LGA, this was an adequate angle for commanding. The DSN was able to sweep the uplink into the LGA and get the telecommunication subsystem coherent while the spacecraft was imaging. Thus, when the spacecraft arrived at Earth-point, the uplink was seamlessly handed from the LGA to the HGA and the spacecraft's downlink remained coherent through the entire process. This technique shaved precious minutes off of the acquisition process and enabled the entire do-si-do concept. Four-minutes was allocated for the DSN to lockup and configure appropriately for the first of the seven do-si-do contacts. Two-minutes were allocated for DSN lockup and configuration for subsequent contacts.

The challenge with this technique was not the short times allocated to acquisition, nor the numerous acquisitions per DSN track, but rather, the station operators were so good and attentive, that during a demonstration of the do-si-do in mid-2010, the station managed to lock onto a side-lobe of the HGA before the slew was complete. When the null of the HGA passed over Earth, the station dropped lock and had to reacquire. This risk was alleviated by having the station operators refrain from locking onto the signal until the time given in the predicts.

In addition to the relative-timed engineering sequence and the three detector sequences, the do-si-do portion also included seven relative-timed playback sequences. Each playback sub-sequence was called from the do-si-do engineering sequence after the appropriate DSN lockup time. A second concern during this portion of Approach was the science team desired more data than can fit in the spacecraft's memory. Thus, some data had to be removed after each do-si-do downlink, before additional data was collected. These deletions were executed during the turns back to comet-pointing.

VIII.C. TCM-22

The navigation team predicted there might be the need for an additional TCM two-days before reaching the comet. Statistically, there was less than a 2-sigma probability this maneuver would be needed. Despite these slim odds, the flight team prepared two additional absolute-timed sequences to support the maneuver planned for 2010-306T15:00:00: one sequence that would kill the Late-Approach backbone and return the spacecraft to the nominal cruise configuration during the seventh, and last do-si-do contact, and a second that continued imaging after TCM-22 completed. This mini-imaging backbone began by slewing the spacecraft from the cruise attitude to the playback attitude. It continued by calling the comet-staring sub-sequences. 11-hours later, the mini-backbone called the do-si-do subsequences that halted the comet-staring subsequences before they could run to completion.

Four-days from the comet, the navigation team decided that they had underestimated the thrust imparted on the comet by the jets of material, seen in Figure 4. In an impressive feat, the navigation team was able to rapidly update their jet model, generate a new orbit estimation of the comet, and design a maneuver profile. TCM-22 was executed using the aforementioned plan flawlessly, restoring the spacecraft back on its planned path towards encounter.

VIII.D. Preparing for Encounter

The final task of Late-Approach was to return the spacecraft back to its nominal cruise configuration to hand off to the Encounter team. The telecommunication subsystem was already appropriately configured. During the seventh, and last contact of the final do-si-do, the spacecraft was commanded to turn back to the nominal cruise attitude while downlinking images. Once image transmission was complete, the backbone sequence deleted all associated sub-sequences and utility sequences and left the spacecraft in its default configuration for the Encounter team.

IX. Conclusion

The Approach phases of the encounter with the comet Hartley-2 were exceedingly successful. During the 59-days of operations, a total of 27,382 unique MRI images, 16,487 unique HRI images, and 20,030 unique infrared spectra were collected by the spacecraft and received uncorrupted on the ground; a total of 63,899 unique images or spectra. These data sets allowed the navigation team to better understand the comet's orbit to target the spacecraft for near perfect flyby geometry and enabled the science team to understand the comet's rotation rate and coma activity. Keep in mind that throughout all of approach, the 2.2 km-long comet was never larger than a single HRI-pixel.

The encounter with Hartley-2 and the EPOXI mission as a whole serves an an excellent example of what can be accomplished with a small, skilled team of scientists and engineers given well-defined constraints. The team was not encumbered with a large bureaucratic infrastructure nor required to prepare for a deluge of reviews. Rather, there was a single main review, with a handful of smaller, peer reviews. The flight team was given modest oversight that did not second guess the team's decisions, as long as they were backed up with a healthy amount of modeling and analysis.

It is the opinion of the author that no spacecraft should be retired. Rather, operations should continue or the spacecraft re-tasked for other uses until it ceases to function. It is difficult to understand how different communities could make use of a spacecraft unless you ask them all.

Acknowledgements

The authors would like to extend their sincerest thanks to everyone on the EPOXI flight team. This is not an inclusive list, but its a start: Brian Kennedy and Matt Abrahamson, AutoNav (JPL); Steve Collins and Mike Luna, ADCS (JPL); Heidi Hallowell, Dustin Putnam, and Meredith Larson, ADCS (Ball); Rich Benson, Alvin Hewitt, and Mark Carmichael, DSN (JPL); Don Hampton, instruments (Univ. of Alaska); Vaughn Cable and Jim Taylor, telecom (JPL); John Valdez and Jeff Lasco, thermal (Ball); Roger Lapthorne, power (Ball); Tomas Ryan, Colleen Jones, and David Holshouser, flight software (Ball); Bob Wing, Lavin Zhang, and Kristine Fong, ground data systems (JPL); Shyam Bhaskaran, Brian Rush, Steve Chesley, Bill Owen, and many others, navigation (JPL); Terry Himes and Eve Pereira, sequencing (JPL); Yee Lee, system administration (JPL); Dennis Wellnitz, Mike A'Hearn (UMD), Ken Klaasan, John Diehl (JPL), and their cohorts, Science; Jeff Pinner, Hal Uffelman, and Genevie Yang, fellow activity leads (JPL); Beth Sholes, propulsion (Ball); Karen Yetter and Jacqueline Jai, DSN Scheduling (JPL); Aimee Meyer, public outreach (JPL); Al Nakata, mission manager (JPL); Tim Larson, project manager (JPL); Beatriz Abu-Ata, project secretary (JPL); Calina Seybold and Jennifer Rocca, contingency planning (JPL); Ben Toyoshima, Bill Currie, and Steve Hillabrand, mission controllers (JPL); Amy Walsh, systems and C&DH (Ball); Leticia Montanez and Don Gibbs, testbed operations, JPL; the station operators at Goldstone, Madrid, and Canberra complexes; and lastly, Steve Wissler, spacecraft team chief (JPL), who's combined knowledge of astronomy, the spacecraft, and the ground system, made the EPOXI mission possible; our single point of failure.

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[1]A'Hearn, M. F., Belton, M. J. S., Delamere, W. a., Kissel, J., Klaasen, K. P., McFadden, L. a., Meech, K. J., Melosh, H. J., Schultz, P. H., Sunshine, J. M., Thomas, P. C., Veverka, J., Yeomans, D. K., Baca, M. W., Busko, I., Crockett, C. J., Collins, S. M., Desnoyer, M., Eberhardy, C. a., Ernst, C. M., Farnham, T. L., Feaga, L., Groussin, O., Hampton, D., Ipatov, S. I., Li, J.-Y., Lindler, D., Lisse, C. M., Mastrodemos, N., Owen, W. M., Richardson, J. E., Wellnitz, D. D., and White, R. L., "Deep Impact: excavating comet Tempel 1." *Science (New York, N.Y.)*, Vol. 310, No. 5746, Oct. 2005, pp. 258–64.

[2]A'Hearn, M. F., Belton, M. J. S., Delamere, W. A., Feaga, L. M., Hampton, D., Kissel, J., Klaasen, K. P., McFadden, L. A., Meech, K. J., Melosh, H. J., Schultz, P. H., Sunshine, J. M., Thomas, P. C., Veverka, J., Wellnitz, D. D., Yeomans, D. K., Besse, S., Bodewits, D., Bowling, T. J., Carcich, B. T., Collins, S. M., Farnham, T. L., Groussin, O., Hermalyn, B., Kelley, M. S., Li, J.-Y., Lindler, D. J., Lisse, C. M., McLaughlin, S. A., Merlin, F., Protopapa, S., Richardson, J. E., and Williams, J. L., "EPOXI at comet Hartley 2." *Science (New York, N.Y.)*, Vol. 332, No. 6036, June 2011, pp. 1396–400.

[3]Rieber, R. R. and Sharrow, R. F., "The contingency of success: Operations for Deep Impact's planet hunt," *Aerospace conference, 2009 IEEE*, Vol. 2831, 2009, pp. 1–10.

[4]Christiansen, J. L., Ballard, S., Charbonneau, D., Deming, D., Holman, M. J., Madhusudhan, N., Seager, S., Wellnitz, D. D., Barry, R. K., Livengood, T. a., Hewagama, T., Hampton, D. L., Lisse, C. M., and A'Hearn, M. F., "System Parameters, Transit Times, and Secondary Eclipse Constraints of the Exoplanet Systems HAT-P-4, TrES-2, TrES-3, and WASP-3 From the NASA EPOXI Mission of Opportunity," *The Astrophysical Journal*, Vol. 726, No. 2, Jan. 2011, pp. 94.

[5]Livengood, T. A., Deming, L. D., A'hearn, M. F., Charbonneau, D., Hewagama, T., Lisse, C. M., McFadden, L. A., Meadows, V. S., Robinson, T. D., Seager, S., and Wellnitz, D. D., "Properties of an Earth-Like Planet Orbiting a Sun-Like Star: Earth Observed by the EPOXI Mission." *Astrobiology*, Vol. 11, No. 9, Nov. 2011, pp. 907–30.

[6]Hampton, D. L., Baer, J. W., Huisjen, M. a., Varner, C. C., Delamere, A., Wellnitz, D. D., A'Hearn, M. F., and Klaasen, K. P., "An Overview of the Instrument Suite for the Deep Impact Mission," *Space Science Reviews*, Vol. 117, No. 1-2, March 2005, pp. 43–93.

[7]Lisse, C., A'Hearn, M., Farnham, T., Fernandez, Y., Groussin, O., Meech, K., and Reach, W., "Recovery and Characterization of the NASA EPOXI Mission Target Comet 85P/Boethin," *Spitzer Proposal ID #460*, Vol. 1, Aug. 2007, pp. 460.

[8]Showstack, R., "In Brief: En route to comet Hartley 2," *Eos, Transactions American Geophysical Union*, Vol. 88, No. 52, Dec. 2007, pp. 579.

[9]Bayer, T. J., "In-Flight Anomalies and Lessons Learned from the Mars Reconnaissance Orbiter Mission," *Aerospace Conference, 2008 IEEE*, March 2008, pp. 1–13.

[10]Rush, B. P., Owen, William M., J., Bhaskaran, S., and Synnott, S. P., "Optical navigation for the EPOXI mission," *21st Annual AAS/AIAA Space Flight Mechanics Meeting*, 2011.